



US 20100121577A1

(19) **United States**
(12) **Patent Application Publication**
Zhang et al.

(10) **Pub. No.: US 2010/0121577 A1**
(43) **Pub. Date: May 13, 2010**

(54) **THREE-DIMENSIONAL LIDAR-BASED CLEAR PATH DETECTION**

Publication Classification

(75) Inventors: **Wende Zhang**, Shelby Township, MI (US); **Kevin Peterson**, Pittsburgh, PA (US)

(51) **Int. Cl.**
G08G 1/16 (2006.01)
H04N 7/18 (2006.01)
H04N 13/02 (2006.01)
G06K 9/00 (2006.01)
G01C 3/08 (2006.01)

Correspondence Address:
CICHOSZ & CICHOSZ, PLLC
129 E. COMMERCE
MILFORD, MI 48381 (US)

(52) **U.S. Cl.** 701/301; 348/143; 348/46; 382/104; 356/4.01; 348/E13.074; 348/E07.085

(73) Assignees: **GM GLOBAL TECHNOLOGY OPERATIONS, INC.**, DETROIT, MI (US); **CARNEGIE MELLON UNIVERSITY**, PITTSBURGH, PA (US)


(57) **ABSTRACT**

(21) Appl. No.: **12/581,861**
(22) Filed: **Oct. 19, 2009**

A method for detecting a clear path of travel for a vehicle includes generating a datastream corresponding to a three-dimensional scan of a target area surrounding the vehicle from a vehicle LIDAR system, estimating a ground plane for a present vehicle location using the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle, and comparing the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle with the estimated ground plane to detect a clear path of vehicle travel.

Related U.S. Application Data

(63) Continuation-in-part of application No. 12/474,594, filed on May 29, 2009, which is a continuation-in-part of application No. 12/108,581, filed on Apr. 24, 2008.






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6A




FIG. 6B




FIG. 6C




FIG. 7




FIG. 8

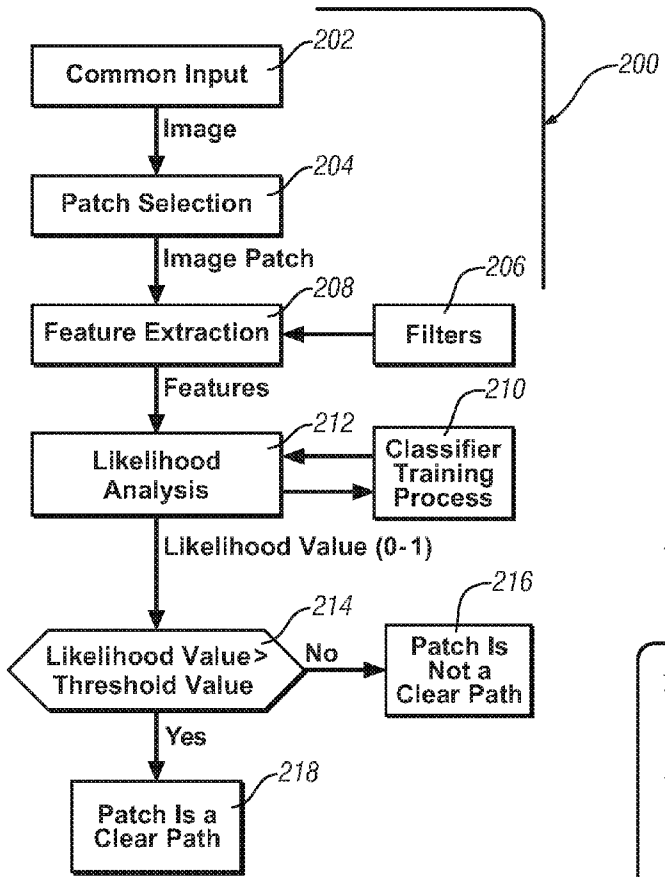


FIG. 9

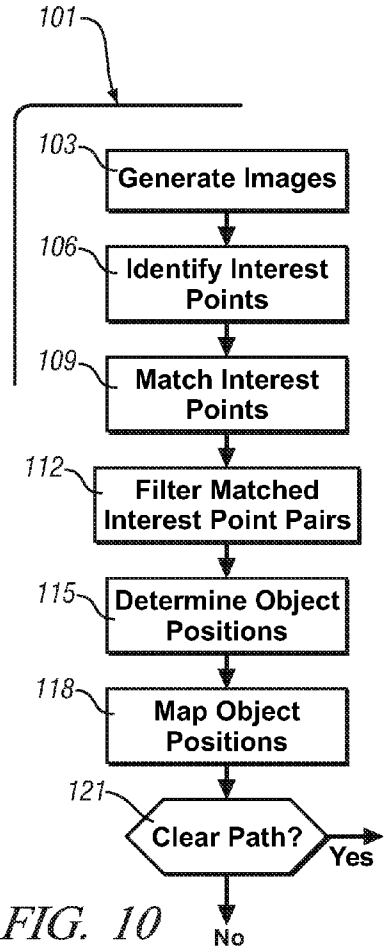


FIG. 10




FIG. 11




FIG. 12




FIG. 13A




FIG. 13B




FIG. 14




FIG. 15

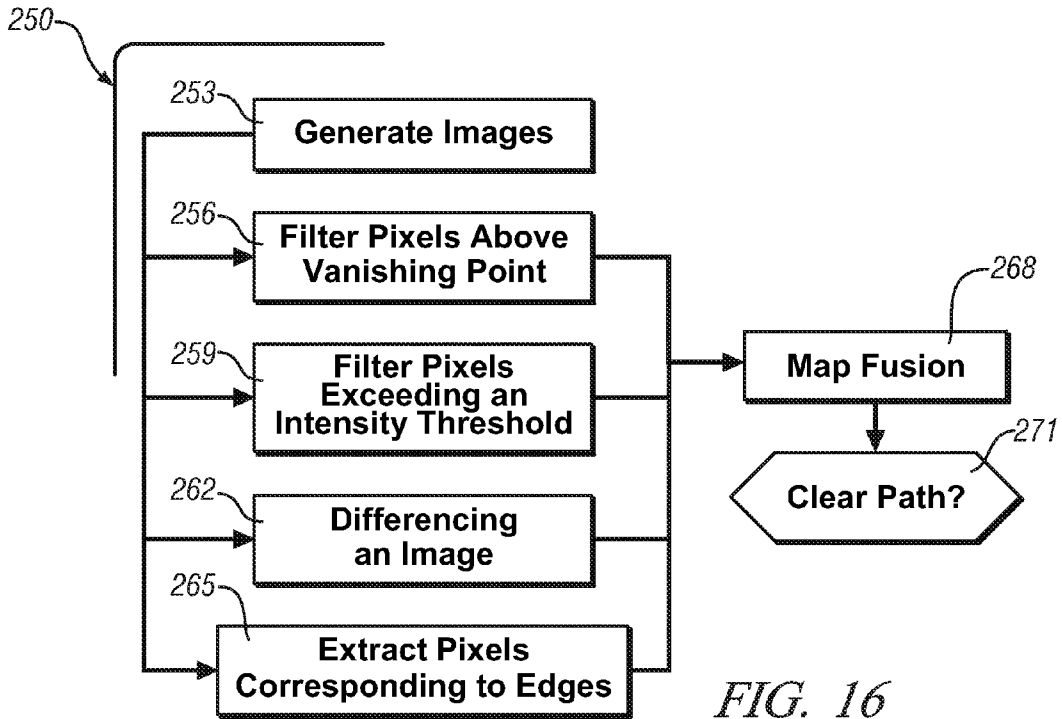


FIG. 16




FIG. 17




FIG. 18







FIG. 21

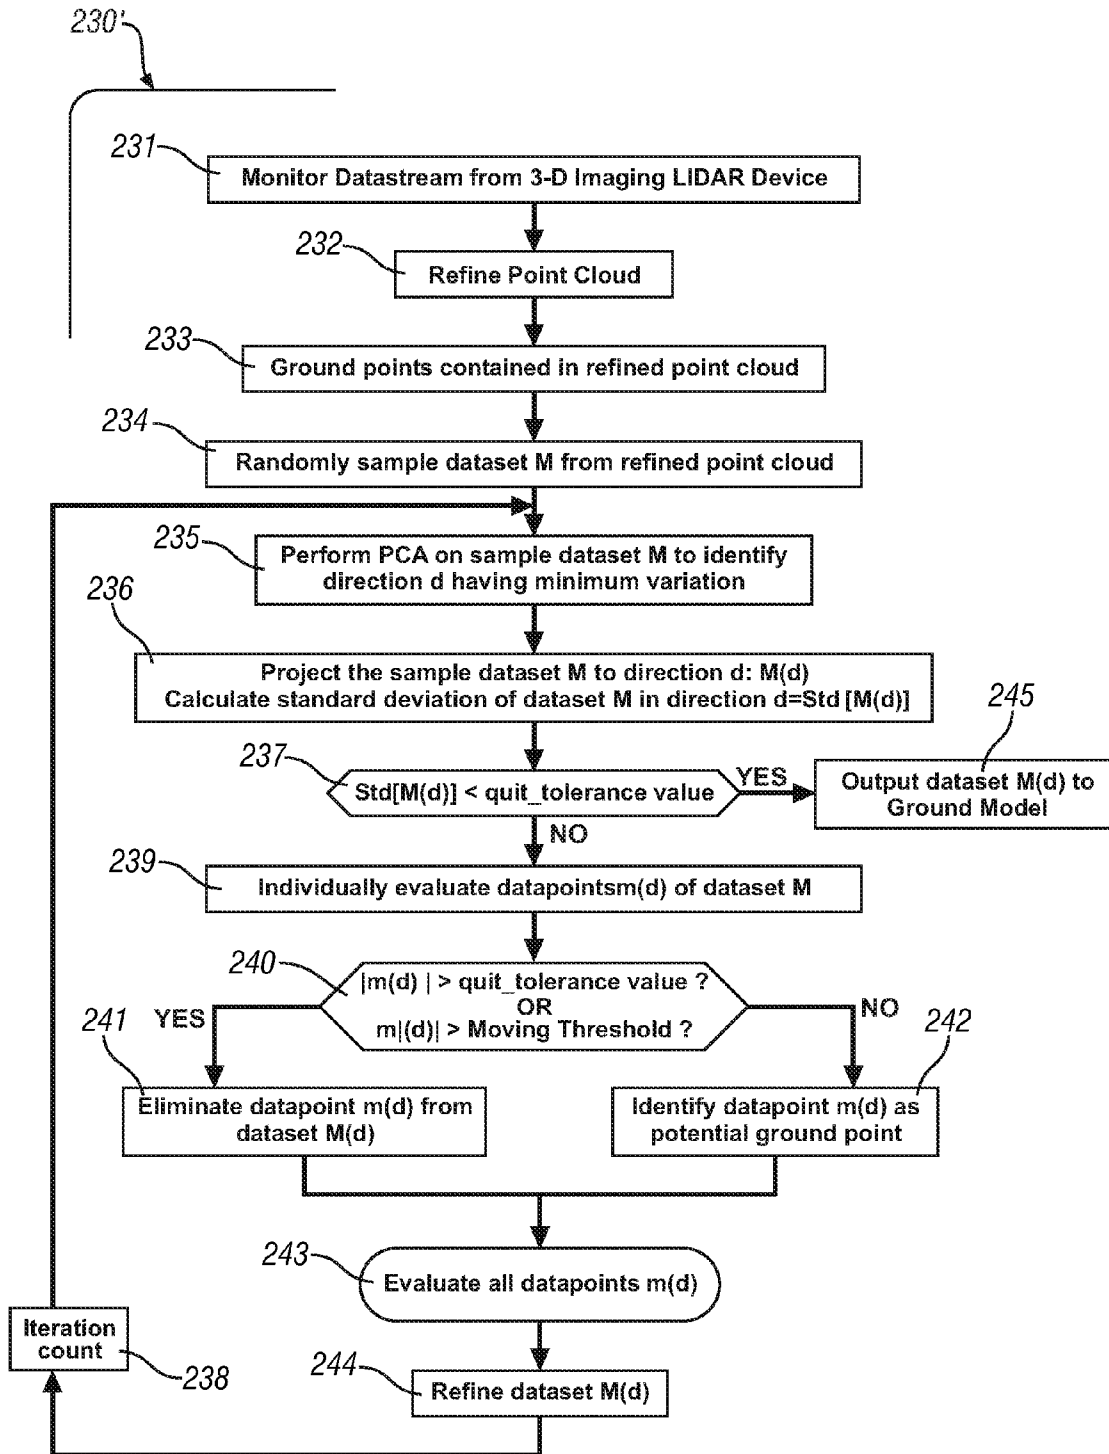


FIG. 22

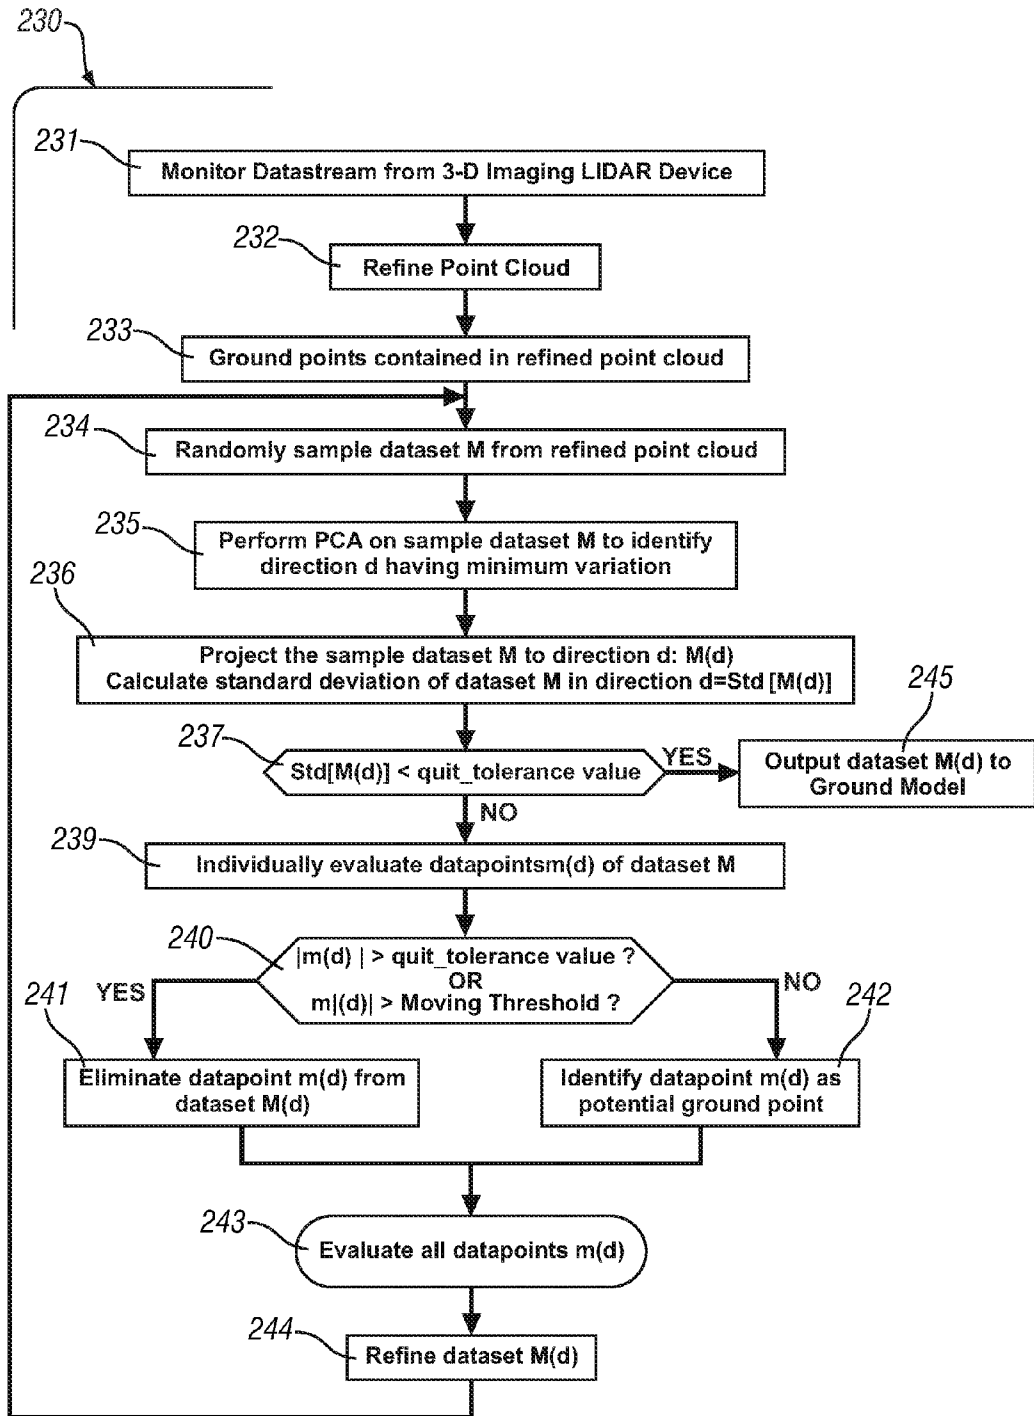


FIG. 23

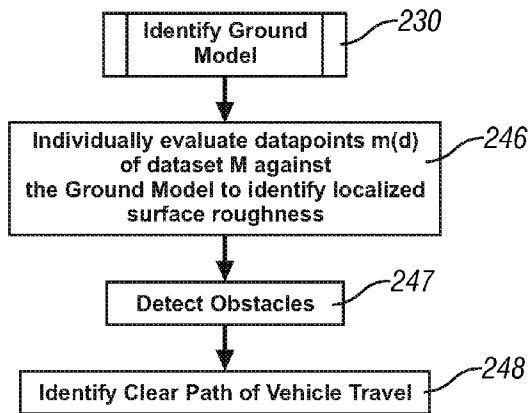


FIG. 24




FIG. 25




FIG. 26

THREE-DIMENSIONAL LIDAR-BASED CLEAR PATH DETECTION

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation-in-part of U.S. application Ser. No. 12/474,594 filed on May 29, 2009, which is a continuation-in-part of U.S. application Ser. No. 12/108,581 filed on Apr. 24, 2008. U.S. application Ser. No. 12/474,594 and U.S. application Ser. No. 12/108,581 are incorporated herein by reference.

TECHNICAL FIELD

[0002] This disclosure is related to automated or semi-automated control of a motor vehicle.

BACKGROUND

[0003] The statements in this section merely provide background information related to the present disclosure and may not constitute prior art.

[0004] Autonomous driving systems and semi-autonomous driving systems utilize inputs regarding the road and other driving conditions to automatically control throttle and steering mechanisms. Accurate estimation and identification of a clear path over which to desirably operate a motor vehicle is critical to successfully replacing the human mind as a control mechanism for vehicle operation.

[0005] Road conditions can be complex. Under normal operation of a vehicle, the human operator makes hundreds of observations per minute and adjusts operation of the vehicle on the basis of perceived road conditions. One aspect of perceiving road conditions is the perception of the road in the context of objects in and around the roadway and navigating a clear path through any objects. Replacing human perception with technology preferentially includes some means to accurately perceive objects and continue to effectively navigate around such objects.

[0006] Technological means for perceiving an object include data from visual cameras and radar imaging. Cameras translate visual images in the form of radiation such as light patterns or infrared signatures into a readable data format. One such data format includes pixelated images, in which a perceived scene is broken down into a series of pixels. Radar imaging utilizes radio waves generated by a transmitter to estimate shapes and objects present in front of the transmitter. Patterns in the waves reflecting off these shapes and objects can be analyzed and the locations of objects can be estimated.

[0007] Once data has been generated regarding the ground in front of the vehicle, the data must be analyzed to estimate the presence of objects from the data. By using cameras and radar imaging systems, ground or roadway in front of the vehicle can be analyzed for the presence of objects that might need to be avoided. However, the mere identification of potential objects to be avoided does not complete the analysis. An important component of any autonomous system includes how potential objects identified in perceived ground data are processed and manipulated to identify a clear path in which to operate the vehicle.

[0008] One known method to identify a clear path in which to operate the vehicle is to catalog and provisionally identify all perceived objects and identify a clear path in light of the locations and behaviors of identified objects. Images may be processed to identify and classify objects according to their

form and relationship to the roadway. While this method can be effective in identifying a clear path, it requires a great deal of processing power, requiring the recognition and separation of different objects in the visual image, for instance, distinguishing between a tree along the side of the road and a pedestrian walking toward the curb. Such methods can be slow or ineffective to process complex situations or may require bulky and expensive equipment to supply the necessary processing capacity.

SUMMARY

[0009] A method for detecting a clear path of travel for a vehicle includes generating a datastream corresponding to a three-dimensional scan of a target area surrounding the vehicle from a vehicle LIDAR system, estimating a ground plane for a present vehicle location using the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle, and comparing the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle with the estimated ground plane to detect a clear path of vehicle travel.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] One or more embodiments will now be described, by way of example, with reference to the accompanying drawings, in which:

[0011] FIG. 1 illustrates an exemplary arrangement of a vehicle equipped with a camera and a radar imaging system in accordance with the disclosure;

[0012] FIG. 2 illustrates a known method to determine a clear path for autonomous driving in accordance with the disclosure;

[0013] FIG. 3 illustrates an exemplary method to determine a clear path utilizing a likelihood analysis of an image in accordance with the disclosure;

[0014] FIG. 4 illustrates an exemplary method to analyze an image in accordance with the disclosure;

[0015] FIG. 5 illustrates an exemplary method to define a classification error by tuning a single threshold in accordance with the disclosure;

[0016] FIGS. 6A, 6B, and 6C illustrate an exemplary determination of an image difference by calculating an absolute image intensity difference in accordance with the disclosure;

[0017] FIG. 7 illustrates an exemplary method to classify a feature as a portion of a clear path and as a detected object at the same time as a method of image analysis in accordance with the disclosure;

[0018] FIG. 8 further illustrates an exemplary method to classify a feature as a portion of a clear path and as a detected object at the same time as a method of image analysis in accordance with the disclosure;

[0019] FIG. 9 illustrates an exemplary process to analyze an image through likelihood analysis in accordance with the disclosure;

[0020] FIG. 10 is a first processing scheme for detecting a clear path, in accordance with the disclosure;

[0021] FIG. 11 shows an exemplary current image including an exemplary set of interest points identified on the current image, in accordance with the present disclosure;

[0022] FIG. 12 shows an exemplary current image captured subsequent to the image depicted in FIG. 11, in accordance with the present disclosure;

[0023] FIGS. 13A and 13B show an exemplary matched point pair, the matched point pair comprising an interest point from a current image and a corresponding interest point from a previous image, and an exemplary pixel region surrounding the points, in accordance with the present disclosure;

[0024] FIG. 13A depicts the interest point from the current image and the pixel region surround the interest point;

[0025] FIG. 13B depicts the interest point from the previous image and the pixel region surround the interest point;

[0026] FIG. 14 graphically illustrates use of an exemplary current image, an exemplary previous image, and operation of a host vehicle to determine the position of a viewed feature, in accordance with the present disclosure;

[0027] FIG. 15 graphically illustrates an exemplary overhead map describing vertical positions of interest points at x and y coordinates in front of the host vehicle, in accordance with the present disclosure;

[0028] FIG. 16 shows a second processing scheme for detecting a clear path, in accordance with the present disclosure;

[0029] FIG. 17 shows an exemplary filtered image based upon pixel color intensity, in accordance with the present disclosure;

[0030] FIG. 18 shows an exemplary filtered image utilizing an edge recognition method, in accordance with the present disclosure;

[0031] FIG. 19 depicts an alternate exemplary process for the second processing scheme, in accordance with the present disclosure;

[0032] FIG. 20 shows an exemplary arrangement of a vehicle equipped with a three-dimensional LIDAR system in accordance with the disclosure;

[0033] FIG. 21 shows an exemplary image from a vehicle-based three-dimensional LIDAR system in accordance with the disclosure;

[0034] FIGS. 22, 23, and 24 schematically show flowcharts in accordance with the disclosure; and

[0035] FIGS. 25 and 26 show exemplary processed images from a vehicle-based three-dimensional LIDAR system in accordance with the disclosure.

DETAILED DESCRIPTION

[0036] Referring now to the drawings, wherein the showings are for the purpose of illustrating certain exemplary embodiments only and not for the purpose of limiting the same, FIG. 1 illustrates an exemplary arrangement of camera 110 located on the front of a vehicle 100 and pointed toward the ground in front of the vehicle 100 in accordance with the disclosure. Camera 110 is in communication with processing module 120 containing logic to process inputs from camera 110. The vehicle 100 may also be equipped with a radar imaging system 130, which, when present, is also in communication with processing module 120. It should be appreciated by those having ordinary skill in the art that the vehicle 100 could utilize a number of methods to identify road conditions in addition or in the alternative to the use of camera 110 and the radar imaging system 130, including GPS information, information from other vehicles in communication with the vehicle 100, historical data concerning the particular roadway, biometric information such as systems reading the visual focus of the driver, a radar imaging system, or other similar systems. The disclosure herein can be applied to various device arrangements and is therefore not limited thereby.

[0037] The camera 110 is a device well known in the art capable of translating visual inputs in the form of light, infrared, or other electro-magnetic (EM) radiation into a data format readily capable of analysis, e.g., a digital, pixelated image. In one embodiment, the camera 110 uses a charge coupled device (CCD) sensor to generate images indicating a field-of-view. Preferably, the camera 110 is configured for continuous image generation, e.g., 30 images generated per second. Images generated by the camera 110 may be stored in memory within the camera 110 or transferred to the processing module 120 for storage and/or analysis. Preferably, each image generated by the camera 110 is a two-dimensional image of known pixel dimensions comprising a plurality of identifiable pixels. The plurality of identifiable pixels may be stored and analyzed using an array. Each pixel may be represented in the array as a set of bits or a plurality of sets of bits wherein the bits correspond to a color on a predetermined palette or color map. Each pixel may be expressed as a function of a plurality of color intensity values such as in a red-green-blue (RGB) color model or a cyan-magenta-yellow-key (CMYK) color model. Preferably, each pixel comprises a plurality of sets of bits wherein each set of bits corresponds to a color intensity and a color intensity value e.g., a first set of bits corresponds to a red color intensity value, a second set of bits corresponds to a green color intensity value, and a third set of bits corresponds to blue color intensity value on the RGB color model.

[0038] The radar imaging device 130 is a device well known in the art incorporating a transmitter capable of emitting radio waves or other EM radiation, a receiver device capable of sensing the emitted waves reflected back to the receiver from objects in front of the transmitter, and means to transfer the sensed waves into a data format capable of analysis, indicating for example range and angle from the objects off which the waves reflected. Alternatively, the radar imaging device 130 may be replaced or supplemented with a light detection and ranging (LIDAR) system configured to transmit and receive optical energy. The received optical energy may be used to determine object geometric dimensions and/or geometrical proximity to the vehicle 100. It will be noted that radar imaging device 130 is optional and unnecessary to perform many of the methods disclosed herein, wherein processing of visual images is capable of accomplishing clear path detection.

[0039] The processing module 120 is illustrated in FIG. 1, and described herein as a discrete element, such illustration is for ease of description and it should be recognized that the functions performed by this element may be combined in one or more devices, e.g., implemented in software, hardware, and/or application-specific integrated circuitry. Preferably, the processing module 120 is a general-purpose digital computer comprising a microprocessor or central processing unit, storage mediums comprising non-volatile memory including read only memory and electrically programmable read only memory, random access memory, a high speed clock, analog to digital and digital to analog circuitry, and input/output circuitry and devices and appropriate signal conditioning and buffer circuitry. The processing module 120 has a set of processing algorithms, comprising resident program instructions and calibrations stored in the non-volatile memory and executed to provide the respective functions of each computer. The algorithms are preferably executed during preset loop cycles. Algorithms are executed by the central processing unit and are operable to monitor inputs from the afore-

mentioned sensing devices and execute control and diagnostic routines to control operation of the actuators, using preset calibrations. Loop cycles may be executed at regular intervals, for example each 3.125, 6.25, 12.5, 25 and 100 milliseconds during ongoing engine and vehicle operation. Alternatively, algorithms may be executed in response to occurrence of an event.

[0040] The processing module 120 executes algorithmic code stored therein to monitor related equipment such as camera 110 and radar imaging system 130 and execute commands or data transfers as indicated by analysis performed within the processing module. Processing module 120 may include algorithms and mechanisms to actuate autonomous driving control by means known in the art and not described herein, or processing module 120 may simply provide information to a separate autonomous driving system. Processing module 120 is adapted to receive input signals from other systems and the operator as necessary depending upon the exact embodiment utilized in conjunction with the control module.

[0041] FIG. 2 illustrates a known method to determine a clear path for autonomous driving in accordance with the disclosure. Image 10 is generated corresponding to the roadway in front of vehicle 100. Through one of various methods, objects 40A, 40B, and 40C are identified within image 10, and each object is categorized and classified according to filtering and trained object behaviors. Separate treatment of each object can be computationally intensive, and requires expensive and bulky equipment to handle the computational load. An algorithm processes all available information regarding the roadway and objects 40 to estimate a clear path available to vehicle 100. Determination of the clear path depends upon the particular classifications and behaviors of the identified objects 40.

[0042] FIG. 3 illustrates an exemplary method to determine a clear path for autonomous or semi-autonomous driving in accordance with the disclosure. Image 10 is depicted including ground 20, horizon 30, and objects 40. Image 10 is collected by camera 110 and represents the road environment in front of vehicle 100. Ground 20 represents the zone of all available paths open to travel without any discrimination on the basis of objects that might be present. The method of FIG. 3 determines a clear path upon ground 20 starts by presuming all of ground 20 is clear, and then utilizes available data to disqualify portions of ground 20 as not clear. In contrast to the method of FIG. 2 which classifies every object 40, the method of FIG. 3 instead analyzes ground 20 and seeks to define a likelihood from available data that some detectable anomaly which may represent object 40 limits or makes not clear that portion of ground 20. This focus upon ground 20 instead of objects 40 avoids the complex computational tasks associated with managing the detection of the objects. Individual classification and tracking of individual objects is unnecessary, as individual objects 40 are simply grouped together as a part of the overall uniform limitation upon ground 20. Ground 20, described above as all paths open to travel without discrimination, minus limits placed on ground 20 by areas found to be not clear, define clear path 50, depicted in FIG. 3 as the area within the dotted lines, or an area with some threshold likelihood of being open for travel of vehicle 100.

[0043] Object 40 that creates not clear limitations upon ground 20 can take many forms. For example, an object 40 can represent a discrete object such as a parked car, a pedestrian, or a road obstacle, or object 40 can also represent a less

discreet change to surface patterns indicating an edge to a road, such as a road-side curb, a grass line, or water covering the roadway. Object 40 can also include an absence of flat road associated with ground 20, for instance, as might be detected with a large hole in the road. Object 40 can also include an indicator without any definable change in height from the road, but with distinct clear path implications for that segment of road, such as a paint pattern on the roadway indicative of a lane marker. The method disclosed herein, by not seeking to identify object 40 but merely to take visual cues from ground 20 and anything in proximity to the ground in image 10, evaluates a likelihood of clear versus not clear and adjusts the control of vehicle 100 for the presence of any object 40.

[0044] Numerous methods for automated analysis of two-dimensional (2D) images are possible. Analysis of image 10 is performed by an algorithm within processing module 120. FIG. 4 illustrates one exemplary method which may be applied to analyze image 10 in accordance with the disclosure. This method subdivides image 10 and identifies a sub-image or patch 60 of ground 20 for analysis, extracts features or analyzes the available visual information from patch 60 to identify any interesting or distinguishing features within the patch, and classifies the patch according to a likelihood of being a clear path according to analysis of the features. Patches with greater than a certain threshold of likeliness are classified as clear, and a compilation of patches can be used to assemble a clear path within the image.

[0045] Patch 60, as a sub-image of image 10, can be identified through any known means, such as random search or swarm search of image 10. Alternatively, information regarding the presence of an object 40 available from some other source of information, such as radar imaging system 130, can be used to identify a patch to analyze the portion of image 10 which should describe object 40. Image 10 may require many patches 60 to analyze the whole image. In addition, multiple overlaying patches or patches of different size could be used to fully analyze a region of image 10 containing information of interest. For instance, a small patch 60 might be used to analyze a small dot on the road; however, a large patch 60 might be required to analyze a series of dots which in isolation might seem uninteresting, but in context of the entire series, could indicate an object 40 of interest. In addition, the resolution of patches applied to a particular area may be modulated based upon information available, for instance, with more patches being applied to a region of image 10 wherein an object 40 is thought to exist. Many schemes or strategies can be utilized to define patches 60 for analysis, and the disclosure is not intended to be limited to the specific embodiments described herein.

[0046] Once a patch 60 has been identified for analysis, processing module 120 processes the patch by application of a filter to extract features from the patch. Additionally, processing module 120 may perform analysis of the location of the patch in context to the location of the vehicle. Filters utilized may take many forms. Filtering algorithms utilized to extract features often search the available visual information for characteristic patterns in the data, defining features by line orientation, line location, color, corner characteristics, other visual attributes, and learned attributes. Learned attributes may be learned by machine learning algorithms within the vehicle, but are most frequently programmed offline and may

be developed experimentally, empirically, predictively, through modeling or other techniques adequate to accurately train distinguishing attributes.

[0047] Once features in patch **60** have been extracted, the patch is classified on the basis of the features to determine the likelihood that the patch is a clear path. Likelihood analysis is a process known in the art by which a likelihood value or a confidence is developed that a particular condition exists. Applied to the present disclosure, classification includes likelihood analysis to determine whether the patch represents a clear path or if ground **20** in this patch is limited by an object **40**. Classification is performed in an exemplary embodiment by application of classifiers or algorithms trained with a database of exemplary road conditions and interactions with detected objects. These classifiers allow processing module **120** to develop a fractional clear path likelihood value for patch **60**, quantifying a confidence between zero and one that the features identified within the patch do not indicate a limiting object **40** which would inhibit free travel of vehicle **100**. A threshold confidence can be set, defining the clear path likelihood required to define the patch as a clear path, for instance by the following logic:

$$\text{Confidence} = \text{ClearPathLikelihood}(i)$$

$$\text{If_Confidence} > 0.5, \text{ then_patch} = \text{clearpath} \quad (1)$$

In this particular exemplary embodiment, a confidence of 50% or 0.5 is selected as the threshold confidence. This number can be developed experimentally, empirically, predictively, through modeling or other techniques adequate to accurately evaluate patches for clear path characteristics.

[0048] The likelihood analysis, as mentioned above, may be performed in one exemplary embodiment by application of trained classifiers to features extracted from a patch. One method analyzes the features a-priori using a training set of images. In this training stage, distinguishing features are selected from a raw feature set, the distinguishing features being defined by methods known in the art, such as Haar wavelet, Gabor wavelet, and Leung-and-Malik filter bank. In addition, 2D image location information based on each feature's minimal classification errors, calculated as the sum of false acceptance rate (FAR) and false rejection rate (FRR), may be utilized by tuning a single threshold as illustrated in FIG. 5. This classification error can be described through the following expression:

$$\text{ClassificationError}(i) = \text{FAR}_i + \text{FRR}_i \quad (2)$$

Information from the trained classifiers is used to classify or weight the feature as indicating a clear path or not clear path, the particular classification depending upon the strength of comparisons to the trained data. Classification of the feature, if the feature is the only feature within the patch, may be directly applied to the patch. Classification of a patch with multiple features identified may take many forms, including the patch being defined by the included feature most indicative of the patch being not clear or the patch being defined by a weighted sum of all of the features included therein.

[0049] The above method can be utilized to examine an individual image **10** and estimate a clear path **50** based upon visual information contained within image **10**. This method may be repeated at some interval as the vehicle travels down the road to take new information into account and extend the formulated clear path to some range in front of the vehicle's new position. Selection of the interval must update image **10** with enough frequency to accurately supply vehicle **100** with

a clear path in which to drive. However, the interval can also be selected to some minimum value to adequately control the vehicle but also not to unduly burden the computational load placed upon processing module **120**.

[0050] Clear path detection can be accomplished through a single image **10** as described above. However, processing speed and accuracy can be improved with the addition of a second image taken in close time proximity to the original image, such as sequential images from a streaming video clip. A second image allows direct comparison to the first and provides for updated information regarding progression of the vehicle and movement of detected objects. Also, the change of perspective of camera **110** allows for different analysis of features from the first image: a feature that may not have shown up clearly or was indistinct in the first image may display at a different camera angle, stand out more distinctly, or may have moved since the first image, allowing the classification algorithm an additional opportunity to define the feature.

[0051] Processing of a second image in relation to the original image **10** can be performed by calculating an image difference. If the image difference of a point of interest, such as a feature identified by radar, is not zero, then the point can be identified as embodying new information. Points where the image difference does equal zero can be eliminated from analysis and computation resources may be conserved. Methods to determine image difference include absolute image intensity difference and vehicle-motion compensated image difference.

[0052] Determining an image difference by calculating an absolute image intensity difference can be used to gather information between two images. One method of absolute image intensity difference includes determining equivalent image characteristics between the original image and the second image in order to compensate for movement in the vehicle between the images, overlaying the images, and noting any significant change in intensity between the images. A comparison between the images indicating a change in image intensity in a certain area contains new information. Areas or patches displaying no change in intensity can be de-emphasized in analysis, whereas areas displaying clear changes in intensity can be focused upon, utilizing aforementioned methods to analyze patches on either or both captured images.

[0053] FIGS. 6A, 6B, and 6C illustrate an exemplary determination of an image difference by calculating an absolute image intensity difference in accordance with the disclosure. FIG. 6A depicts an original image. FIG. 6B depicts a second image with changes from the original image. In particular the depicted circular shape has shifted to the left. A comparison of the two images as illustrated in FIG. 6C, an output representing the result of an absolute image intensity difference comparison, identifies one region having gotten darker from the first image to the second image and another region having gotten lighter from the first image to the second image. Such a method can be described as differencing. Analysis of the comparison yields information that some change as a result of movement or change of perspective is likely available in that region of the images. In this way, absolute image intensity difference can be used to analyze a pair of sequential images to identify a potentially not clear path.

[0054] Likewise, determining an image difference by calculating a vehicle-motion compensated image difference can be used to gather information between two images. Many methods to calculate a vehicle-motion compensated image

difference are known. One exemplary method of vehicle-motion compensated image difference includes analyzing a potential object as both a stationary portion of a clear path and a detected object at the same time. Likelihood analysis is performed on features identified corresponding to the potential object from both classifications at the same time, and the classifications may be compared, for example, through the following logic:

$$\text{Confidence}(i) = \text{ClearPathLikelihood}(i) - \text{DetectedObjectLikelihood}(i)$$

$$\text{If_Confidence} > 0, \text{ then_patch} = \text{clearpath} \quad (3)$$

In this exemplary comparison, if confidence(i) is greater than zero, then the patch containing the feature is classified as a clear path. If confidence(i) equals or is less than zero, then the patch containing the feature is classified as not a clear path or limited. However, different values may be selected for the confidence level to classify the patch as a clear path. For example, testing may show that false positives are more likely than false negatives, so some factor or offset can be introduced.

[0055] FIG. 7 illustrates one method to classify a feature as a portion of a clear path and as a detected object at the same time as described above in accordance with the disclosure. Image 10 includes object 40, trapezoidal projection 70, and rectangular projection 80. This method utilizes an assumption projecting object 40 as a flat object on the ground within projection 70 to test the classification of the feature as a portion of a clear path. The method also utilized an assumption projecting object 40 as a vertical object within rectangular projection 80 to test the classification of the feature as a detected object. FIG. 8 illustrates comparisons made in data collected between the two images to evaluate the nature of object 40 in accordance with the disclosure. Camera 110 at time t₁ observes and captures data from object 40 in the form of a first image. If object 40 is an actual detected object, the profile observed by camera 110 of object 40 at time t₁ will correspond to point 90A. If object 40 is a flat object in the same plane as ground 20, then the profile observed by camera 110 of object 40 at time t₁ will correspond to point 90B. Between times t₁ and t₂, camera 110 travels some distance. A second image is captured at time t₂, and information regarding object 40 can be tested by applying an algorithm looking at visible attributes of the object in the second image in comparison to the first image. If object 40 is an actual detected object, extending upward from ground 20, then the profile of object 40 at time t₂ will be observed at point 90C. If object 40 is a flat object in the same plane as ground 20, then the profile of object 40 at time t₂ will be observed at point 90B. The comparison derived through vehicle-motion compensated image difference can directly assign a confidence by application of classifiers based on the observations of points 90, or the comparison may simply point to the area displaying change as a point of interest. Testing of the object against both classifications, as a flat object and as an actual detected object, allows either the area including object 40 to be identified for further analysis through analysis of a patch as described above or direct development of a clear path likelihood and a detected object likelihood for comparison, for example as in logic expression (3) above.

[0056] Information available from analysis of the second image can additionally be improved by integration of information regarding movement of the vehicle, such as speed and yaw-rate. Information regarding vehicle motion is available

from a number of sources, including the vehicle speedometer, vehicle dynamic sensors or wheel speed sensors, anti-lock braking mechanisms, and GPS location systems. Algorithms may utilize this vehicle movement information, for example, in conjunction with the projections described in FIGS. 7 and 8 to project angles which should exist in a feature laying flat on the ground in the second image based upon data from the first image and the measured movement of the vehicle between the images.

[0057] The number of images utilized for comparison need not be limited to two. Multiple image analysis can be performed at multiple iterations, with an object being tracked and compared over a number of cycles. As mentioned above, computational efficiency can be gained by utilizing image difference analysis to identify points of interest and eliminating areas with zero difference from subsequent analyses. Such efficiencies can be used in multiple iterations, for example, saying that only points of interest identified between a first and second image will be analyzed in the third and fourth images taken. At some point, a fresh set of images will need to be compared to ensure that none of the areas showing zero difference have had any change, for example a moving object impinging upon a previously clear path. The utilization of image difference analyses and of focused analyses, eliminating areas identified with zero change, will vary from application to application and may vary between different operating conditions, such as vehicle speed or perceived operating environment. The particular utilization of image difference analyses and of focused analyses can take many different embodiments, and the disclosure is not intended to be limited to the specific embodiments described herein.

[0058] FIG. 9 illustrates an exemplary process 200 wherein input from a camera is analyzed to determine a clear path likelihood in accordance with the disclosure. Camera input in the form of an image is generated at step 202. At step 204, a patch is selected for analysis from the image. Step 206 represents a filter or set of filters available to process the patch. At step 208, feature extraction is performed upon the selected patch through application of filters available from step 206 and application of other algorithms. Step 210 includes a classifier training process. As mentioned above, classifiers or logic used in developing likelihood values are initially trained offline. Training may optionally be continued in the vehicle based upon fuzzy logic, neural networks, or other learning mechanisms known in the art. These trained classifiers are utilized in step 212 to perform a likelihood analysis upon the features extracted in step 208, and a likelihood value for the patch is developed. This likelihood value expresses a confidence that the selected patch is clear. At step 214, the likelihood value developed in step 212 is compared to a threshold likelihood value. If the likelihood value is greater than the threshold value, then at step 218 the patch is identified as a clear path. If the likelihood value is not greater than the threshold value, then the patch is identified as a not clear path. As described above, process 200 may be repeated or reiterated in a number of ways, with the same image being analyzed repeatedly with the selection and analysis of different patches, and an identified patch can be tracked and analyzed for change over a number of sequential images.

[0059] As mentioned above, processing module 120 may include algorithms and mechanisms to actuate autonomous driving control by means known in the art and not described herein, or processing module 120 may simply provide information to a separate autonomous driving system. Reactions to

perceived objects can vary, and include but are not limited to steering changes, throttle changes, braking responses, and warning and relinquishing control of the vehicle to the operator.

[0060] The above method, utilizing analysis of patches, is one method to establish a clear path of travel for a vehicle. Two additional, related processing schemes are disclosed, employing similar methods to analyze pixels instead of patches. A first processing scheme utilizing texture-rich methods to analyze images to identify a clear path is disclosed. FIG. 10 shows a first processing scheme 101 for detecting a clear path, described herein as comprising discrete elements. Such illustration is for ease of description and it should be recognized that the functions performed by these elements may be combined in one or more devices, e.g., implemented in software, hardware, and/or application-specific integrated circuitry.

[0061] Processing scheme 101 is an exemplary method to apply texture-rich image analysis of a field-of-view in front of the host vehicle and begins with step 103, wherein images of the field-of-view in front of the vehicle are generated. At step 106, the monitored images are examined for interest points, for example, examining pixel color intensity as described hereinabove and comparing the pixel or a group of pixels with surrounding pixels. Through methods known in the art, interest points can be identified and utilized to accomplish methods described herein. At step 109, sequential images of the field-of-view in front of the vehicle are compared when the vehicle is in motion, and interest points from each image are matched to corresponding points in sequential images which correspond to the same points in the field-of-view, where possible. Matching includes locating corresponding points through template matching or comparing interest points on the sequential images, taking into account movement of the host vehicle, and making a best estimate whether two points represent the same object or feature visible in the field-of-view. While interest points can be matched, not all matched corresponding point pairs represent high quality corresponding point pairs that allow the identification of their three-dimensional positions in the field-of-view for classifications as a clear path for the vehicle to travel through. At step 112, a filter is applied to the matched corresponding point pairs in order to identify high quality corresponding point pairs that can be utilized for three-dimensional position identification with high confidence. At step 115, the high quality corresponding point pairs are analyzed to determine three-dimensional positions of objects represented by the corresponding point pairs. It will be appreciated that corresponding points at different heights as compared to ground level will move differently between sequential images. Analyzing movement of the interest points can yield estimated three dimensional coordinates of the interest points. At step 118, the determined object positions are utilized to map object positions in front of the host vehicle. At step 121, the map is utilized to determine a clear path in front of the host vehicle.

[0062] FIG. 11 shows an exemplary current image (k) 500 including an exemplary set of interest points 501 identified on the current image (k), each interest point 501 corresponding to a pixel. The processing module 120 preferably identifies interest points 501 on each image generated including the current image (k) 500 and a previous image (k-1). An interest point 501 is an identifiable pixel on an image and is associated with a set of visual information, i.e., rich-texture features, and is associated with objects located in the field-of-view. The

exemplary field-of-view or view depicted in FIG. 11 includes a road surface 510, sections in the road surface 515, curbs 520 and 521, lane markers 522, intersecting lane 526, and buildings 524 and 525. Candidate interest points 501 are identified in the view by interest point extraction programming from which a set of interest points 501 may be selected. The interest point extraction programming is executed on the images by one of several known methods, e.g., a scale-invariant feature transform (SIFT), methods employing corner detection or other shape detection, or a Sobel filter. The interest point extraction programming is preferably executed in the processing module 120, but may be combined in one or more devices, e.g., implemented in software, hardware, and/or application-specific integrated circuitry. The interest point extraction programming locates candidate interest points 501 in each image corresponding to predetermined identifiable rich-texture features, e.g., pixels indicating an edge, pixels indicating a transition in the visual data, wherein potentially significant features can be identified in the view. In the exemplary view of FIG. 11, of the many identified interest points 501 identified, one point, 501A, is identified at a corner formed at a section 505 in curb 520.

[0063] After identifying candidate interest points 501, the interest point extraction programming may filter the candidate interest points by removing redundant candidate interest points, i.e., candidate interest points corresponding to a same feature. For example, the interest point extraction programming may filter multiple candidate interest points corresponding to an edge, resulting in less candidate interest points corresponding to the edge. The remaining set of candidate interest points are the interest points for the image. In one embodiment, filtering redundant candidate interest points is executed for computational efficiency.

[0064] Once the interest points 501 in a current image (k) 500 are identified, the processing module 120 matches the set of interest points from the current image (k) with a set of points identified from the previous image (k-1) to determine a set of matched point pairs, corresponding to step 109 in process 101.

[0065] A first method to determine a set of matched point pairs includes using correspondence matching programming to match interest points from a current image (k) and a previous image (k-1). The processing module 120 matches the set of interest points from the current image (k) with a set of interest points identified in the previous image (k-1) to determine a set of matched interest point pairs. FIG. 12 shows an exemplary current image (k) captured subsequent to the image depicted in FIG. 11. The exemplary current image (k) 530 represents an image captured sequentially to the image depicted in FIG. 11, wherein the perspective of the viewer has progressed slightly along the viewed road between the images. In this exemplary situation, image 500 of FIG. 11, although a current image (k) at the time it was captured, now can act as a previous image (k-1) to current image (k) 530. The exemplary view depicted in FIG. 12 includes road surface 510, sections in the road surface 515, curbs 520 and 521, lane markers 522, and buildings 524 and 525, as depicted in FIG. 11, with slightly different perspectives of every feature in the view. In the exemplary view of FIG. 12, many identified interest points 531 are identified according to methods described above, and one point, 531A, is identified at a corner formed at a section 505 in curb 520. It will be noted that interest points 531 identified in image 530 are not necessarily directly corresponding to interest points 501 identified in

image 500. Points 531 identified in image 530 are processed and filtered as described in association with points 501, as described in association with FIG. 11. Preferably, several interest points 531 from the set of interest points from the current image (k) 530 are matched with an interest point 501 from the set of interest points identified from the previous image (k-1) 500 to determine a plurality of matched point pairs. Each interest point comprising a matched point pair is expected to correspond to a same feature associated with an object in the view. To determine the set of matched point pairs, the set of interest points 531 from the current image (k) are compared to the set of interest points 501 identified from the previous image (k-1) by one of several known correspondence matching programming, e.g., a scale-invariant feature transform (SIFT) feature matching programming and optical flow programming. The correspondence matching programming is preferably executed in the processing module 120, but may be combined in one or more devices, e.g., implemented in software, hardware, and/or application-specific integrated circuitry. The resulting matched point pairs correspond to a same feature located on both the current image (k) and the previous image (k-1) wherein the same feature is associated with a same object in the view.

[0066] A second method to determine a set of matched point pairs includes using template matching programming to match interest points from the current image (k) and the previous image (k-1). In one embodiment templates are generated using a predetermined pixel region neighboring an interest point. An exemplary region 535 is depicted in association with interest point 531A in FIG. 12. The template matching may be determined using one of multiple methods, including one of several known template matching programming methods to find the corresponding points of interest in the previous image, e.g., Lucas-Kanade or Horn-Schunck. Preferably, the templates neighboring an interest point in a current image and compared with templates neighboring an interest point in a previous image. When the template matching programming determines that the templates match, the interest points are included in the set of matched point pairs. The resulting matched point pairs correspond to a same feature located on both the current image (k) and the previous image (k-1) wherein the same feature is associated with a same object in the view.

[0067] FIGS. 13A and 13B show an exemplary matched point pair, the matched point pair comprising interest point 531A from the current image (k) 530 and a second interest point 501A from the previous image (k-1) 500, and an exemplary pixel region surrounding the points. According to the method described above, pixel region 535 around interest point 531A is depicted in FIG. 13A. As described above, interest point 531A is a point identified at a corner of section 505 between portions of curb 520. Pixel region 535 is selected around point 531A and preferably includes discernable features, textures, or patterns that can be used to positively identify the region. FIG. 13B similarly depicts interest point 501A and a pixel region 540 around the interest point. Because of the change in perspective between image 500 and image 530, it is possible that pixel region 535 and pixel region 540 can include some differences, but applying methods described above, a comparison of the pixel regions and identifying features therein can be utilized as one method to confirm with some confidence that interest points 531A and 501A represent the same feature within the view captured in images 500 and 530 and may be treated as a matched pair.

[0068] After the matched point pairs are determined, the processing module 120, by applying filters removing low quality matched corresponding point pairs indicative of poorly matched or mismatched pairs, selects high quality, preferential matched point pairs from the set of matched point pairs, as described in exemplary step 112 above. Preferential matched point pairs may be selected based upon quality control criteria. In one embodiment, matched point pairs are analyzed and identified as preferential matched point pairs after satisfying each criterion.

[0069] A first criterion is met when a distance between points of a matched point pair is less than a threshold. The distance is determined based upon the location of the points as if they were located on a same two-dimensional image coordinates. The threshold may be predetermined and location-dependent or dynamic based upon speed of the vehicle. A second criterion is met when both points of a matched point pair are a predetermined threshold distance from an image boundary. Due to the perspective of the camera to objects at the edges of the image boundary and vehicle motion, points too close to the edge either cannot locate a corresponding point, which may be located outside of the perspective of the camera or can create distorted estimates of object locations. Additionally, under normal vehicle operation in a forward direction, the probability of a vehicle needing information regarding a clear path toward the extremes of the view is small. A third criterion is met when a difference in color between a predetermined area neighboring each point of a matched point pair is less than a threshold difference in color. Utilizing the exemplary pixel regions defined in FIG. 13, a difference in color between the pixel regions 535 and 540 can be used to increase confidence in the matched pair. Each pixel within each of the regions can be assigned a number value based upon its corresponding color. The assigned number may be based upon a bit value determined while capturing the image or may be based upon an indexed color referencing a predetermined palette. The processing module 120 calculates absolute differences between the assigned number values within the first area 501 and the second area 502 pixel-by-pixel, and sums the differences, which is the difference in color. The difference is compared with the threshold difference. If the difference is less than the threshold difference the third criterion is met. The threshold difference in color can be selected by any method sufficient to calibrate accurate road presence or clear path estimation. Different threshold differences in color can be utilized based upon lighting levels exterior to the vehicle, weather, speed of the vehicle, or any other factors that affect accurate estimation of clear path presence based upon differences in color. By judging whether the areas around the interest points are similar, a judgment can be made regarding whether the same area including the feature represented by the interest points is being analyzed in both images.

[0070] It will be appreciated that the three criteria described above are exemplary criteria useful to judge matched pairs, but some portion of these criteria or additional unnamed, similar criteria can be utilized to judge validity of matched pairs. Selection of criteria to judge matched pairs can additionally be made based upon conditions such as light level present outside the vehicle, weather, speed of the vehicle, and any other factor affecting an ability to judge matched pairs or an urgency to quickly and accurately define a clear path.

[0071] After selecting the preferential matched point pairs, the processing module 120 determines positions of a plurality

of points associated with features of objects in the view with respect to the vehicle **100**, as described in exemplary step **115** above. FIG. **14** graphically illustrates use of an exemplary current image, an exemplary previous image, and operation of a host vehicle to determine the position of a viewed feature. Object position in a horizontal frame of reference and a height of the object as compared to a ground level can be determined based upon the preferential set of matched point pairs within sequential images **317** and **327** (the matched pair demonstrated in this figure through points **1** and **2**), a distance (d) the vehicle **100** traveled from a first position **310** to a second position **320**, and vehicle yaw (θ). The processing module **120** executes one of several known triangulation methods to determine the position of the point with respect to the vehicle **100** and a height of the point. In FIG. **14**, a field-of-view **315** from **310** including an image $k-1$ **317** is depicted, and a field-of-view **325** from **320** including an image k **327** is depicted. Distance d is depicted describing a distance that an observer has traveled between points **310** and **320**. Lines longitudinally oriented to fields of view **315** and **325** are represented by lines **312** and **322**, respectively. An angular change in orientation or yaw (θ) is depicted describing an angle between lines **312** and **322**. Distance d can be determined an exemplary vehicle by tracking speed of the vehicle through a sample time between the time at **310** and the time at **320**. Similarly, θ can be determined by tracking a yaw rate of the vehicle through the sample time. Object being viewed **305** is depicted, and lines **318** and **328** are depicted describing lines of view from points **310** and **320** to object **305**, respectively. Points **1** and **2** are depicted upon images **317** and **327** where lines **318** and **328** intersect images **317** and **327**, respectively. Distance (a) can be defined describing a location of point **1** upon image **317**, and distance (b) can be defined describing a location of point **2** upon image **327**. It will be appreciated that FIG. **14** represents an overhead view wherein a position of object **305** can be defined in a horizontal plane, and similar computations in a side view of the same objects can be utilized to define a vertical position of object **305** with relation a known ground level for the observer. By applying well known methods of triangulation, positional data such as distances a and b from sequential images and vehicular data such as vehicle speed, vehicle yaw rate, and sample time can be utilized to locate a position of a viewed object in relation to the vehicle and determine relative motion of the object in relation to the vehicle. These methods of triangulation can yield a position of the object in a horizontal plane and a height of the object in relation to a ground level.

[**0072**] Once the position and height are determined, the processing module **120** may plot the points on an overhead map, as described in exemplary step **118** above. FIG. **15** graphically illustrates an exemplary overhead map describing vertical positions of interest points at x and y coordinates in front of the host vehicle. Position $\langle 0, 0 \rangle$ on the x axis and y axis corresponds to the second position **320** of the vehicle **100** described hereinabove, or the current position of vehicle **100**. Preferably, object height is classified into predetermined categories. For example, objects with minimal height, e.g., below a predetermined threshold, may be classified as ground, objects exceeding ground height but less than a second predetermined threshold may be classified into a second category which is close to and preferably less than a vehicle height, and objects greater than the second predetermined threshold may be classified into a third category. As FIG. **15** shows, objects with minimal height are classified as ground (Ground), objects exceeding ground height but less than a 2

meter threshold are classified into the second category (Less than 2 m), and objects greater than the 2 meter threshold are classified into the third category (Greater than 2 m).

[**0073**] After plotting the objects on the map, the processing module **120** detects a clear path based upon features plotted in the map, as described in exemplary step **121** above. The term “clear path” as used herein is to be given its ordinary and customary meaning to a person of ordinary skill in the art (and it is not to be limited to a special or customized meaning), and refers without limitation to a path free of objects exceeding a threshold, e.g., a path free of objects classified in the above-mentioned second and third categories. Features of objects that are classified in predetermined categories, e.g., the second and third category described hereinabove are identified by the processing module **120** as not-clear areas, i.e., not desirable for vehicle travel. Preferably, a predetermined area neighboring each object classified in the predetermined categories is identified by the processing module **120** as not-clear area. The processing module **120** may determine that a clear path is present on the map using one of multiple methods. A first method for determining a clear path comprises comparing an anticipated path with the map including the not-clear areas. If the anticipated path does not intersect any not-clear areas then the processing module **120** determines that the anticipated path is a clear path. However, if the anticipated path intersects a not-clear area then the processing module **120** determines that a clear path is not present. A second method for determining a clear path comprises using the not-clear areas in the map to determine a clear path. Any path not intersecting the not-clear areas may be used as a clear path.

[**0074**] The above method utilizes sequential images to establish a map of object positions and vertical heights in front of the vehicle, such that a clear path can be defined. It will be appreciated that in any two given images, a particular object might not be classified as including two high quality interest points sufficient to be mapped in that particular analysis. However, the above analysis occurs multiple times per second of vehicle travel. As the vehicle travels forward through the clear path, different perspectives upon an object will be gained and a large number of images will be analyzed. Travel over a path and analysis of the multiple iterative images through that path build a confidence through the analyses that no object contradicting the clear path exists in the identified clear path.

[**0075**] FIG. **16** shows a second processing scheme for detecting a clear path. Whereas the first processing scheme described herein utilizes texture rich methods, analyzing pixelated features that describe distinct interest points based upon a contextual view within the image, the second disclosed processing scheme can be described as an exemplary texture-less method of image analysis, filtering from an image non-conforming regions of the image as not belonging to a planar, consistent road surface. By filtering from an image non-conforming regions, a clear surface can be identified from the remaining image as a potential clear path for the vehicle to travel. Process **250** is illustrated in FIG. **16**, and described herein as comprising discrete elements. Such illustration is for ease of description and it should be recognized that the functions performed by these elements may be combined in one or more devices, e.g., implemented in software, hardware, and/or application-specific integrated circuitry. For example, the process **250** may be executed as one or more algorithms in the processing module **120**. The process **250** for detecting a clear path comprises creating a plurality of filtered images and fusing the filtered images together to determine a clear path.

[0076] Exemplary process 250 is disclosed, including steps to employ an exemplary texture-less method of clear path detection. Process 250 begins at step 253, wherein images are generated. Steps 256, 259, 262, and 265 describe alternate exemplary methods to filter an image to identify a location of a clear path including a road surface. It will be appreciated that each of the methods processes the image in one way to help identify the clear path. A process could be employed with some portion of the four exemplary methods or a process could be employed to include unnamed but similar methods to process the image. Any method can be employed that filters within an image a clear surface upon which a road surface can be estimated from other portions of the image that do not indicate a potential road surface. Step 256 applies an exemplary method to filter an area below a horizon or vanishing point, including a road surface that can be driven upon, from an area above the horizon or vanishing point, including sky and other vertical features that cannot be part of a road surface. Step 259 applies a filter based upon variance in pixel intensity, based upon a premise that a road surface will include a large surface with a visual intensity very common across the surface. Step 262 applies a filter based upon differencing sequential images, allowing analysis of changes between the images. Step 265 applies a filter based upon identifying pixels representing edges or transitions in the visual data. Applying the various methods in parallel, the results can be fused into a single map of the image in step 268 and analyzed for visual data indicative of a clear path of travel in step 271.

[0077] A first filtered image is created using a vanishing point, as is described in exemplary step 256 above. The term “vanishing point” as used herein is a broad term, and is to be given its ordinary and customary meaning to one ordinarily skilled in the art, and refers to an infinite far point on the horizon that is intersected by multiple parallel lines on the ground in the view. Identifying a road surface creating a clear path on which to drive is necessarily below the vanishing point or horizon line. Filtering images to only analyze an area below the horizon line helps to clarify the pixels being analyzed to identify a road surface from irrelevant pixels. As one skilled in the art appreciates, there are many known methods for determining a vanishing point and corresponding horizon line. One known method comprises determining the vanishing point based upon a point where the lane markers on the left and right of a host vehicle intersect each other. The horizon line is determined based upon the vanishing point. The first filtered image consists of pixels on the current image below the horizon line.

[0078] The processing module 120 creates a second filtered image based upon pixel color intensity of the plurality of pixels comprising the current image (k), as is described in exemplary step 259 above. FIG. 17 shows an exemplary second filtered image based upon pixel color intensity. The second filtered image comprises pixels selected from the plurality of pixels. A number of methods to filter pixels are envisioned. The processing module 120 compares each color intensity value of a pixel with an associated color distribution mean e.g., a red, green, and blue color distribution mean. The color distribution means are an average color intensity distribution for an associated color intensity associated with pixels previously identified in clear paths. The color distribution means may additionally be based upon historically gathered color distribution means of clear path pixels of the associated colors. A pixel is selected for the second filtered image when each color intensity value of the pixel is less than a color intensity threshold from the associated color distribution mean. When one color intensity value of the pixel is greater

than a color intensity threshold from the associated color distribution mean, the pixel is excluded from the second filtered image. The color distribution mean for the associated color intensity is an average color intensity value distributed in the image. The color distribution means are initially predetermined. In one embodiment, the processing module 120 adapts the color distribution mean for each associated color intensity based upon the color intensity distributions in a predetermined number of images.

[0079] As described hereinabove, each pixel comprises a plurality of color intensity values. Preferably, the color intensity value is compared with the associating color distribution mean. For example, for a pixel created using the RGB color model, a first bit set associated with a red color intensity is compared with the color distribution mean associated with the red color intensity, the bit set associated with the green color intensity is compared with the color distribution mean associated with the green color intensity, and the bit set associated with the blue color intensity is compared with the color distribution mean associated with the blue color intensity. If the difference between each color intensity comprising the pixel and the associated color distribution mean is less than a color intensity threshold, then the pixel is used for the second filtered image.

[0080] Using the RGB color model, each pixel comprises a first, second, and third color intensity value. A first color intensity is a red intensity, a second color intensity is a green intensity, and a third color intensity is blue intensity. The pixel for the second filtered image is selected from the plurality of pixels comprising the current image (k) when the difference between a red intensity of the pixel and a red color distribution mean is less than a red intensity threshold, and the difference between a green intensity of the pixel and a green color distribution mean is less than a green intensity threshold, and the difference between a blue intensity of the pixel and a blue color distribution mean is less than a blue intensity threshold. The following equation expresses this process using the RGB color model:

$$\begin{aligned} |R-R_M| < \text{threshold}_R \text{ and } |G-G_M| < \text{threshold}_G \text{ and } \\ |B-B_M| < \text{threshold}_B \end{aligned} \tag{4}$$

wherein

- [0081] R is a red color
- [0082] G is a green color,
- [0083] B is a blue color,
- [0084] R_M is a color distribution mean for the red color,
- [0085] G_M is a color distribution mean for the green color,
- [0086] B_M is a color distribution mean for the blue color,
- [0087] threshold_R is a red intensity threshold,
- [0088] threshold_G is a green intensity threshold, and
- [0089] threshold_B is a blue intensity threshold.

[0090] One skilled in the art will appreciate that Equation 4 above may be configured to be used with multiple other color models including the CMYK, YUV, and LAB color models.

[0091] The processing module 120 creates a third filtered image using the current image (k) and the previous image (k-1) by differencing the images, as is described in exemplary step 262 above. In an exemplary process to create a differenced image, a third filtered image is created as an image difference of the current image (k) and the previous image (k-1) and comprises a pixel by pixel comparison representing the difference between the current image (k) and the previous image (k-1). Such a process is described above in association with FIGS. 6A-6C. A first method for determining the third filtered image comprises determining a difference between individual pixels on a current image (k) and corresponding individual pixels on the previous image (k-1). The difference may be determined by subtracting color inten-

sity values associated with a pixel from corresponding color intensity values to determine a color intensity difference, e.g., subtracting the red color intensity value of a first pixel on the current image (k) from the red color intensity value of a first pixel on the previous image (k-1). The resulting color intensity difference comprises the pixel value and corresponds to the difference between the pixel on the current image (k) and the previous image (k-1). Each pixel value may be stored in an array corresponding to the third filtered image. In one embodiment, an absolute value of the difference is determined and stored in the array. The resulting array comprises values representing the difference between the two images and comprises the image difference. After determining the array, the processing module 120 removes values representing pixels that changed less than the predetermined threshold. The resulting values can, in one embodiment, be replaced with a predetermined value, e.g., one, to represent on a resulting image a clear map of identified differences between the two utilized images.

[0092] A second exemplary method for determining the third filtered image is based upon differencing images is described. A determination can be made by comparing the current image (k) and a motion adjusted previous image (k-1), and differences between these images can be used to map objects not fitting a particular model used to generate the motion adjusted image. One exemplary model to generate the motion adjusted image includes utilizing triangulation methods, such as the methods describing in association with FIG. 14, information regarding the motion of the vehicle, and an assumption that all detected points are at ground level to predict motion of objects in the previous image (k-1). By using an assumption that all objects in the image are at ground level, predictions of motion of the objects in the resulting image are consistent with the previous image existing as an entirely flat image existing at ground level. Comparison of this resulting motion adjusted previous image (k-1) with the actual current image (k), including perspective changes upon objects not at ground level, allows for identification of all objects or features in the viewed images that are not on the ground level. By subtracting pixels in the compared images, for example according to the first exemplary process for determining the third filtered image described above, pixels representing objects not on the ground level can be identified by the non-zero values or values exceeding a determinative threshold. In this way, the processing module 120 can determine the third filtered image as described hereinabove using the current image (k) and the motion adjusted previous image instead of the previous image (k-1).

[0093] The processing module 120 creates a fourth filtered image based upon color intensity values of the plurality of pixels comprising the current image (k), as described in exemplary step 265 above. FIG. 18 shows a fourth filtered image utilizing an edge recognition method. The fourth filtered image comprises edges of the current image (k). To create the fourth filtered image, the processing module 120 extracts pixels from the image based upon color intensity values that correspond to edges using one of several known edge detection filters, e.g., a Sobel filter. The edge detection filter is preferably executed in the processing module 120, but may be combined in one or more devices, e.g., implemented in software, hardware, and/or application-specific integrated circuitry. In one embodiment, each pixel is analyzed using a Sobel operator. The Sobel operator computes a gradient vector of color intensity at each pixel resulting in a direction of the largest possible increase from light to dark and the rate of change in that direction. Points corresponding to a rate of change exceeding a threshold and corresponding to gradient

vectors at nearby pixels indicate edges and are included in the fourth filtered image. Those pixels may be replaced with a predetermined pixel color, e.g., black.

[0094] After creating the plurality of filtered images, the processing module 120 fuses the filtered images together to determine pixel locations corresponding to a clear path, as described in exemplary step 268. The fused filtered image comprises an image containing pixels contained in all of the plurality of filtered images, i.e., only pixels contained at a particular pixel location in each of the plurality of filtered images are included in the fused filtered image. Fusing the filtered images together comprises determining pixel locations included on each of the plurality of filtered images. Each pixel location is compared with corresponding pixel locations on each of the filtered images. When each of the plurality of filtered images contains a pixel at a corresponding pixel location, that pixel is included on a fused clear path map representing a clear surface on the road. The fused clear path is used to detect a clear path for the vehicle 100. Pixels on the fused clear path map correspond to desirable driving locations in the view. Locations on the fused clear path map without pixels correspond to undesirable driving locations in the view.

[0095] Process 250, described in FIG. 16, applies various methods to images in parallel to identify features in a view relevant to defining a clear path. However, it will be appreciated that these methods need not be performed in parallel, but rather the methods can be used to process images sequentially in steps to identify features in a view relevant to defining a clear path. An exemplary process 600 for the second processing scheme is depicted in FIG. 19. Process 600 initiates by generating a sequence of images at step 603. Each image is delivered in parallel to two steps, and the results from the two steps are fused and made available for subsequent steps in the process. The parallel steps include step 606, wherein a first filtered image according to filtering pixels by the vanishing point method described above in order to filter from processing the portions of the images that cannot represent a road surface, and step 609, wherein a second filtered image by filtering pixels according to intensity, as described above, is utilized to filter from processing portions of the images describing features inconsistent with a road surface. A fused image of the first filtered image and the second filtered image is created in step 612, the fused image including only a portion of the image below the vanishing point and including pixels with intensity consistent with a road surface. Sequential fused images are used, in step 615, to create a third filtered image using one of the differencing methods described above. In step 618, the third filtered image is utilized to create a fourth filtered image using the edge extraction methods described above. The results of process 600 can then be analyzed for visual data indicative of a clear path of travel.

[0096] Processes 250 and 600 illustrate two exemplary arrangements in which filtering methods can be utilized in combination to process a sequence of images to identify a clear path. However, it will be appreciated that such combinations can take many forms, steps can be applied in different orders, and fewer or additional filters can be used in different combinations. Additionally, filters can be utilized selectively or configurations of processes can be utilized selectively, for example, with different combinations being utilized in the daytime and others being utilized in the night. For example, some combinations can continue to analyze road surfaces illuminated by headlights whereas other methods might not be useful without an entirely illuminated view. In another example, different combinations can be utilized when rain or snow is present on the road. For example, some methods and

analysis methods might continue to be useful with a light coating of snow, for example, identifying tire paths in the snow as potential clear paths, whereas other methods might not be useful where a white covering hides most identifiable features. In another example, different combinations can be utilized when additional information is made available, for example, with infrared, radar, or GPS data augmenting or facilitating filtering of the images based upon the additional information. A number of combinations of filtering methods are envisioned, and the disclosure is not intended to be limited to the particular embodiments described herein.

[0097] As described above, a first processing scheme, employing texture-rich methods, and a second processing scheme, employing texture-less methods, are described, each scheme enabling identification of road features useful to describe a clear path in the view of the vehicle. However, it will be appreciated that a single analysis based upon either scheme can yield ambiguous results, for example, wherein particular lighting conditions, shadows from another vehicle, heat caused distortion of the road surface, or other similar issues can cause portions of a road surface to be misidentified as not conforming to the rest of the road surface. One method to resolve ambiguous results is accomplished by analyzing iterative images as the vehicle travels along the road. It will be appreciated that as a vehicle travels along a road, hundreds of images taken in quick succession can be analyzed of a particular section of road. As the vehicle travels forward, different lighting and exposure of the road surface as a result of different viewing perspectives will be observed. Travel over a path and analysis of the multiple iterative images through that path can build confidence through the analyses that the particular portion of the road surface is correctly estimated as a clear path.

[0098] Another method to resolve ambiguous results of either scheme is to utilize both schemes and fuse the results, utilizing the combination of the two schemes to increase confidence of identification of a clear path. Each of the schemes has advantages over the other scheme in different aspects of clear path detection. For example, the exemplary texture rich method excels at detecting objects distinctly standing up from the ground level with rich textures. This method positively identifies objects detectable at different heights and yields a map of areas that a vehicle should not travel lest the objects collide with the vehicle. In another example, the exemplary texture-less method excels at identifying a region wherein pixels identify a commonly appearing surface. This method positively identifies a region wherein a road surface is likely to exist and maps the boundaries of this surface.

[0099] The first scheme and the second scheme can be fused in a number of ways. An image with identified points and determined heights identified with texture-rich methods can be overlaid with a filtered image generated by texture-less methods, and agreement of the two methods can be used to define a clear path through the overlaid image. In an alternative method to fuse the two schemes, the data from each of the two schemes can be used to project information upon a programmed overhead map of an area in front of the vehicle, and this overhead map including data gained from analysis of the two schemes can include building confidence indications for regions of the map. In an alternative method to fuse the two schemes, one scheme can be utilized as a primary or dominant scheme, and the second scheme can be utilized or activated to analyze regions in the view identified as ambiguous or unclear. In any method to fuse the two schemes, strengths of one processing scheme can be used to reduce weaknesses of the other processing scheme. If both schemes concur that the

path is clear, then the processing module employing the schemes may determine with increased confidence that the path is desirable for the vehicle to traverse. A number of methods to fuse the identified schemes are envisioned, and the disclosure is not intended to be limited to the particular embodiments described herein. Additionally, either scheme or both schemes can be combined with the method employed above utilizing analysis of patches.

[0100] One skilled in the art will recognize that lighting normalization may be applied to the captured image data. Normalization is a process which changes the range of pixel intensity values. The purpose of the normalization is to bring the image into a range that is more suitable for machine process to improve the reliability. For example, each pixel value is normalized to be a zero mean and unit variance to enhance the image contrast, specifically in a low lighting environment or when contrast is poor due to glare.

[0101] The filters and algorithms described hereinabove may take many forms. Filtering algorithms utilized to extract features often search the available visual information for characteristic patterns in the data, defining features by line orientation, line location, color, corner characteristics, other visual attributes, and learned attributes. Attributes may be developed experimentally, empirically, predictively, through modeling or other techniques adequate to accurately train distinguishing attributes. Learned attributes may be learned by machine learning algorithms or fuzzy logic within the vehicle adapting over time. Additionally, learned attributes or learned landmarks can be gathered from repetitive travel of a vehicle over a route and utilized in clear path identification.

[0102] FIG. 20 schematically shows an embodiment of a vehicle system **100'** including a three-dimensional (3D) imaging light detection and ranging (LIDAR) optical sensing device system **111**, an optional camera **110** and an optional radar system **130**, and a processing module **120'**. The three-dimensional (3D) imaging light detection and ranging (LIDAR) optical sensing device system **111** is preferably mounted on a top location of the vehicle **100'** to generate LIDAR images corresponding to 360° of rotation about the vehicle **100'**. The 3D imaging LIDAR system **111** emits laser pulses and monitors elapsed time between transmission of each pulse and detection of a reflected signal to determine a range or distance, referred to as a phase-sensitive measurement. The emitted pulses are oriented with reference to an orthogonal 3D reference grid **113** having a locus **113A** fixed on the vehicle **100'** and described in terms of an X-axis, a Y-axis, and a Z-axis that are preferably associated with the vehicle **100'**. Preferably, the X-axis is associated with a direction of travel of the vehicle **100'**, the Y-axis is orthogonal to the direction of travel of the vehicle **100'** when travelling on a horizontal surface, and the Z-axis is an elevation axis that is orthogonal to the X and Y axes. The 3D imaging LIDAR system **111** generates an array of data that is descriptive of a target area surrounding the vehicle **100'** that preferably includes a ground plane for the vehicle **100'** and vertical projections therefrom. The 3D imaging LIDAR system **111** transmits a datastream containing the array of data to the processing module **120'** at a rate of 3 Mbits/sec. The datastream comprises a three-dimensional point cloud, comprising a three-dimensional scan of a target area surrounding the vehicle **100'**, an example of which is depicted with reference to FIG. 21. The three-dimensional point cloud is preferably an array of 3D datapoints generated by the 3D imaging LIDAR system **111** and projecting outwardly from with a single location of the vehicle **100'**, and thus associated with a single location of the locus **113A** during vehicle operation.

[0103] FIG. 22 schematically shows a flowchart 230 for sampling and processing the datastream generated by the 3D imaging LIDAR system 111 comprising the three-dimensional point cloud, preferably executed in the processing module 120' as algorithmic code. Overall, the flowchart 230 depicts processing of the datastream generated by the 3D imaging LIDAR system 111 to identify a ground model comprising a data-based construct associated with a ground plane on which the vehicle 100' is travelling.

[0104] The ground model is used to identify obstacles that can be in the form of vertical projections from the ground plane, as shown with reference to data depicted in FIGS. 25 and 26. Obstacles can be associated with a potential path of travel for the vehicle 100', adjacent to the potential path of travel for the vehicle 100', and other locations. A clear path of travel for the vehicle 100' is identified, comprising a path of travel for the vehicle 100' that avoids the identified obstacles projecting from the identified ground plane.

[0105] In operation, the three-dimensional point cloud generated by the 3D imaging LIDAR system 111 is ongoingly monitored (231). The 3D reference grid 113 having locus 113A is oriented relative to the vehicle 100'. As such, when the vehicle 100' is on a level surface, the 3D reference grid 113 corresponds to a surface grid having x and y axes that are level, and when the vehicle 100' is on a tilted surface, the 3D reference grid 113 does not correspond to a surface grid having x and y axes that are level. Thus, as the vehicle 100' approaches an inclined surface, e.g., a hill, a portion of the three-dimensional point cloud includes a portion of a ground model that has zero tilt, and portion of a ground model that has an incline. Furthermore, during ongoing vehicle travel, orientation of the 3D imaging LIDAR system 111 and the data contained in the datastream is affected by vehicle yaw, pitch, and roll associated with vehicle maneuvering, e.g., cornering and braking, and road conditions including potholes and other irregular surface features.

[0106] Signal processing includes refining the three-dimensional point cloud by identifying and eliminating datapoints from the three-dimensional point cloud that are distal from the locus 113A of the vehicle 100', including any datapoint that has a large Z-axis component or a Y-axis component that is a distance greater than a predetermined threshold associated with vehicle speed and other factors that eliminate the datapoint from being in a potential path of travel for the vehicle 100'. This analysis takes into account the portion of the ground model that has zero tilt, and any portion of the ground model having an incline to refine the three-dimensional point cloud (232). Ground points are contained in the refined three-dimensional point cloud (233).

[0107] A sample dataset M is selected from the refined three-dimensional point cloud. The sample dataset M is a subset of the three-dimensional point cloud and contains a plurality of datapoints m randomly selected using known sampling techniques (234).

[0108] The sample dataset M is subjected to principal component analysis (PCA) techniques to identify a direction d having a minimum variation (235).

[0109] The principal component analysis (PCA) techniques are applied to the datapoints m associated with the sample dataset M, which results in a dataset M(d). Dataset M(d) includes the datapoints m projecting in the d direction, with each of the datapoints includes represented as m(d). The principal component analysis (PCA) techniques include calculating least-square error terms that further result in calculating a standard deviation of the m(d) datapoints associated with the dataset M(d), represented as Std[M(d)] (236).

[0110] The standard deviation Std[M(d)] is compared to a predetermined threshold (quit_tolerance value), which is a calibratable fixed value indicative of a ground model associated with the ground plane (237).

[0111] When the standard deviation Std[M(d)] is less than the predetermined threshold (quit_tolerance value), the sample dataset M(d) is output to the ground model, which is the data-based construct associated with the ground plane on which the vehicle 100' is travelling. The ground model is stored in the processing module 120' as an estimate of the ground plane (245).

[0112] When the standard deviation Std[M(d)] is greater than the predetermined threshold (quit_tolerance value), each individual datapoint m(d) of the dataset M(d) is evaluated (239). Evaluating each individual datapoint m(d) of the dataset M(d) includes comparing the absolute value of each of the individual datapoints projected in the d direction |m(d)| with the predetermined threshold (quit_tolerance value), and with a moving threshold (240). When the absolute value of the datapoint projected in the d direction |m(d)| exceeds either one of the predetermined threshold (quit_tolerance value) and the moving threshold, the datapoint m(d) is eliminated from the M(d) dataset (241). Otherwise, the datapoint m(d) is identified as a potential ground point and retained in the M(d) dataset (242). The process is repeated for each individual datapoint m(d) of the dataset M(d) (243). The dataset M(d) is refined by eliminating those datapoints m(d) that exceed either one of the predetermined threshold (quit_tolerance value) and the moving threshold (244).

[0113] In the embodiment shown, the process iteratively executes to evaluate the dataset M using the updated refined point cloud for a predetermined quantity of iterations, in order to refine the dataset M(d) which is eventually output to the ground model (245). The moving threshold against which the datapoints m(d) are compared is recalculated during each iteration in which the dataset M is evaluated. The dataset M(d) is preferably evaluated and refined multiple times, as indicated by the iteration loop and an associated iteration count (238).

[0114] In one embodiment, the moving threshold can be determined as a linear or other function of the standard deviation Std[M(d)], with the moving threshold decreasing with increased iterations for the dataset M(d). This is performed when the processing speed is sufficient to allow multiple iterations to refine the ground model using the present dataset M(d). Multiple iterations of analyzing the dataset M(d) allows a more refined definition of the ground model, including accounting for surface roughness of a road. Road surface roughness includes accommodating randomly occurring features on the road surface including snow, and various and varying road surfaces including dirt, gravel, grass, asphalt and concrete.

[0115] FIG. 23 schematically shows flowchart 230', wherein the process 230 described with reference to FIG. 22 is repeated during each iteration with an updated dataset M. The sampled portion of the datastream, i.e., each dataset M, is subjected to principal component analysis (PCA) to identify a minimum variation direction d (235).

[0116] Principal component analysis is an analytical method for processing data that includes using an orthogonal linear transformation to transform data, e.g., the dataset M, to a new coordinate system wherein the greatest variance by any projection of data lies on a first coordinate or principal component, the second greatest variance by any projection of data lies on a second coordinate or principal component, and the least variance by any projection of data lies on a third coordinate or principal component for three-dimensional

datapoints. The first, second, and third coordinates for a multivariate system having three dimensional datapoints are referred to as eigenvectors, with the projections of the individual datapoints m that are so associated referred to as projection coefficients. The coordinate system resulting from principal component analysis identifies the first and second eigenvectors corresponding to the greatest variance by projection of the data lying on the first coordinate or principal component and the second greatest variance by projection of the data lying on the second coordinate or principal component. The first and second eigenvectors can be used to define the ground model described herein.

[0117] The third eigenvector is defined as that element having the least variance by the projection of data lying on a third coordinate or principal component. The third eigenvector is that portion of each datapoint that is projected in the d direction as described herein. Thus, each datapoint has a projection coefficient corresponding to the third eigenvector that is represented as $m(d)$. Thus, given a set of datapoints defined in a Euclidean space, each datapoint defined in terms of coordinates associated with the X , Y , and Z axes, the principal component of interest is that which corresponds to the third, least eigenvalue having a line that passes through a mean value and minimizes a sum-squared error for those points. Thus when there is a three-dimensional dataset, the transformed dataset can become a transformed three-dimensional dataset. The principal component analysis permits transformation of the dataset M to eigenvectors. Each datapoint m has projection coefficients corresponding to the first, second, and third eigenvectors, including the projection coefficient $m(d)$ projecting in the third coordinate or the principal component system. The principal component analysis permits transformation of each datapoint m having a location defined in the X , Y , and Z axes to include a transformed datapoint having the projection coefficient $m(d)$, wherein the projection coefficient $m(d)$ for with each datapoint is associated with the coordinate or principal component having the least variance by any projection of data.

[0118] In processing the datastream generated by the 3D imaging LIDAR system **111** comprising the three-dimensional point cloud, there is a need to identify the ground plane to distinguish between ground points that make up the ground plane and obstacles to identify a clear path of travel, in order to construct the ground model using the sampled datapoints in the datastream. Characteristics of the ground plane include that it is a substantially flat surface in a large area, and a substantial portion of the datapoints in the three-dimensional point cloud will be associated with the ground plane. Thus, the application of principal component analysis allows the ground plane to be identified by associating specific ones of the datapoints $m(d)$ with the ground plane, and identifying specific other ones of the datapoints $m(d)$ that are not in the ground plane as obstacles.

[0119] FIG. **24** schematically shows a flowchart for identifying a clear path of travel using the datastream generated by the 3D imaging LIDAR system **111**. This includes identifying the ground model (**230**), the process of which is described with reference to one of FIGS. **22** and **23**. Individual datapoints $m(d)$ in the d direction of the original randomly sampled dataset M (at **234**) are compared to the ground model (**246**). The process includes analyzing and classifying individual datapoints $m(d)$ to determine variances that can be associated with localized surface roughness and smoothness, including subjecting the individual datapoints $m(d)$ of the dataset M to the aforementioned texture-rich analysis processes. This process allows for distinguishing between obstacles and uneven features of the surface of the ground

model. Thus specific datapoints $m(d)$ can be identified as one of a surface feature in the ground model and an obstacle. An individual datapoint $m(d)$ is considered an obstacle when its magnitude exceeds a predetermined threshold that preferably differs from the moving threshold, taking into account the variances that are associated with localized surface roughness. Obstacles can be detected thereby (**247**). Once any obstacles are identified, a clear path of vehicle travel can be identified, comprising that which avoids the identified obstacles.

[0120] In an embodiment using the camera **110** in conjunction with the 3D imaging LIDAR system **111**, obstacles can be identified using information from the 3D imaging LIDAR system **111** and the camera **110**, and a clear path of vehicle travel can be identified, comprising that which avoids the obstacles identified by both systems.

[0121] FIG. **25** graphically shows the three-dimensional point cloud wherein specific ones of the datapoints $m(d)$ are classified as part of the ground plane, and specific other ones of the datapoints $m(d)$ are classified as obstacles.

[0122] FIG. **26** graphically shows the three-dimensional point cloud wherein the datapoints $m(d)$ have been classified as part of the ground plane or as an obstacle, and further refined to include as obstacles only those data points that presently obstruct the travel path of the vehicle, indicating presence of a clear path of travel.

[0123] The disclosure has described certain preferred embodiments and modifications thereto. Further modifications and alterations may occur to others upon reading and understanding the specification. Therefore, it is intended that the disclosure not be limited to the particular embodiment(s) disclosed as the best mode contemplated for carrying out this disclosure, but that the disclosure will include all embodiments falling within the scope of the appended claims.

1. Method for detecting a clear path of travel for a vehicle, comprising

generating a datastream corresponding to a three-dimensional scan of a target area surrounding the vehicle from a vehicle LIDAR system;

estimating a ground plane for a present vehicle location using the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle; and

comparing the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle with the estimated ground plane to detect a clear path of vehicle travel.

2. The method of claim **1**, wherein estimating the ground plane for the present vehicle location comprises executing a principal component analysis on the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle.

3. The method of claim **2**, wherein executing the principal component analysis on the datastream comprises executing the principal component analysis on a dataset sampled from the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle.

4. The method of claim **3**, wherein the principal component analysis on the sampled dataset is iteratively executed.

5. The method of claim **4**, comprising comparing the sampled dataset with the estimated ground plane to detect the clear path of vehicle travel.

6. The method of claim **3**:

wherein executing the principal component analysis on the sampled dataset comprises transforming the sampled dataset into first, second, and third eigenvectors, the

third eigenvector having a least variance associated therewith compared to the first and second eigenvectors; wherein estimating the ground plane comprises estimating the ground plane to correspond to datapoints from the transformed sampled dataset having projection coefficients corresponding to the third eigenvector that are less than a threshold value; and further comprising identifying obstacles comprising datapoints of the transformed sampled dataset having projection coefficients corresponding to the third eigenvector that exceed a second threshold value.

7. The method of claim **1**, further comprising: generating data from a vehicle camera imaging system; and

wherein comparing the datastream further comprises comparing the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle with the data generated by the vehicle camera imaging system to detect the clear path of vehicle travel.

8. Method for detecting a clear path of travel for a vehicle, comprising:

generating a datastream corresponding to a scan of a target area surrounding the vehicle from a three-dimensional vehicle LIDAR system;

estimating a ground plane for a present vehicle location using the datastream corresponding to the scan of the target area surrounding the vehicle; and

detecting a clear path of vehicle travel by comparing the datastream corresponding to the scan of the target area surrounding the vehicle with the estimated ground plane.

9. The method of claim **8**, wherein estimating the ground plane for the present vehicle location comprises executing a principal component analysis on the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle.

10. The method of claim **9**, wherein executing the principal component analysis on the datastream comprises executing the principal component analysis on a dataset sampled from the datastream corresponding to the scan of the target area surrounding the vehicle.

11. The method of claim **10**, wherein the principal component analysis on the sampled dataset is iteratively executed.

12. The method of claim **11**, comprising comparing the sampled dataset with the estimated ground plane to detect the clear path of vehicle travel.

13. The method of claim **10**, further comprising: wherein executing the principal component analysis on the sampled dataset comprises transforming the sampled dataset into first, second, and third eigenvectors, the

third eigenvector having a least variance associated therewith compared to the first and second eigenvectors; wherein estimating the ground plane comprises estimating the ground plane to correspond to datapoints from the transformed sampled dataset having projection coefficients corresponding to the third eigenvector that are less than a threshold value; and

further comprising identifying obstacles comprising datapoints of the transformed sampled dataset having projection coefficients corresponding to the third eigenvector that exceed a second threshold value.

14. The method of claim **8**, further comprising: generating data from a vehicle camera imaging system; and

wherein comparing the datastream further comprises comparing the datastream corresponding to the scan of the target area surrounding the vehicle with the data generated by the vehicle camera imaging system to detect the clear path of vehicle travel.

15. Method for detecting a clear path of travel for a vehicle equipped with a LIDAR system configured to generate a datastream corresponding to a three-dimensional scan of a target area surrounding the vehicle, comprising:

executing a principal component analysis on the datastream corresponding to the three-dimensional scan of the target area surrounding the vehicle;

estimating a ground plane for a present vehicle location using the principal component analysis;

identifying obstacles using the principal component analysis and the estimated ground plane for the present vehicle location; and

detecting a clear path of vehicle travel associated with the estimated ground plane and the identified obstacles.

16. The method of claim **15**:

wherein executing a principal component analysis on the datastream comprises executing the principal component analysis on a sampled dataset of the datastream, including transforming the sampled dataset into first, second, and third eigenvectors, the third eigenvector having a least variance associated therewith compared to the first and second eigenvectors; and

further comprising associating the estimated ground plane with datapoints from the transformed dataset having projection coefficients corresponding to the third eigenvector that are less than a threshold value; and

identifying obstacles comprising datapoints of the transformed sampled dataset having projection coefficients corresponding to the third eigenvector that exceed a second threshold value.

* * * * *